

Amendments to the Specification

Please replace paragraph [0020] with the following amended paragraph:

[0020] Another known approach employs a hosted architecture. A virtual machine application uses a VM driver to load a virtual machine monitor at a privileged level. Typical of this approach are the disclosures of U.S. Patent Nos. 6,075,938 and 6,496,847, which are incorporated herein by reference. The virtual machine monitor then uses the I/O services of a host operating system to accommodate user-level VM applications. Current examples of this approach include the VMware Workstation™, the VMware GSX Server™, both available from VMware, Inc., 3145 Porter Drive, Palo Alto, CA 94304, and the Connectix Virtual PC™, available from Microsoft Corporation, One Microsoft Way, Redmond, WA 98052-6399. Another example is the open source Plex86 Virtual Machine, available via the Internet ~~at the URL~~ <http://plex86.sourceforge.net/>. The hosted architecture is attractive due to its simplicity. However, it incurs a performance penalty because the virtual machine monitor must itself run as a scheduled application under the host operating system, and could even be swapped out. Furthermore, it requires emulators to be written and maintained for diverse I/O devices that are invoked by the virtual machine monitor.

Please replace paragraph [0023] with the following amended paragraph:

[0023] An unsuccessful attempt to implement a VM computing paradigm on cluster-based systems is disclosed in the document *The Memory and Communication Subsystem of Virtual Machines for Cluster Computing*, Shiliang Hu and Xidong Wang, Jan. 2002 (Hu et al.), published on the Internet ~~at the URL~~ <http://www.es.wisc.edu/~wxd/report/eee902.pdf>. In this proposed arrangement, multiple SMP clusters of NUMA-like processors are monitored by virtual machine monitors. A cluster interconnect deals with message passing among the clusters. The system consists of multiple virtual machines that operate under a single operating system, and support parallel programming models. While a virtual computer built according to this paradigm would initially appear to be highly scalable, preliminary simulations of the communication and memory subsystems were discouraging. A further difficulty is posed by limitations of current operating systems, which are generally unaware of the locality of NUMA-type memory. According to Hu et al., the proposed paradigm could not be reduced to practice until substantial technological changes occur in the industry. Thus Hu et al. appears to have encountered a well-known difficulty: cluster machines generally, and NUMA machines in particular, can be scaled up successfully only if some way is found to ensure a high computation to communication ratio in regard to both data distribution and explicit communication among the clusters and processors.